arXiv:1605.01219v1 [q-bio.GN] 4 May 2016

Evolution of transcription factor families along the human lineage.

Antonio Rosanova,[1,*] Matteo Osella,[1] and Michele Caselle[1]

[1]*Department of Physics and INFN, Università degli Studi di Torino, via P.Giuria 1, I-10125 Turin, Italy*

Transcription factors (TFs) exert their regulatory action by binding to DNA with specific sequence preferences. However, different TFs can partially share their binding sequences. This "redundancy" of binding defines a way of organizing TFs in "motif families" that goes beyond the usual classification based on protein structural similarities. Since the TF binding preferences ultimately define the target genes, the motif family organization entails information about the structure of transcriptional regulation as it has been shaped by evolution. Focusing on the human lineage, we show that a one-parameter evolutionary model of the Birth-Death-Innovation type can explain the empirical repartition of TFs in motif families, thus identifying the relevant evolutionary forces at its origin. More importantly, the model allows to pinpoint few deviations in human from the neutral scenario it assumes: three over-expanded families corresponding to HOX and FOX type genes, a set of "singleton" TFs for which duplication seems to be selected against, and an higher-than-average rate of diversification of the binding preferences of TFs with a Zinc Finger DNA binding domain. Finally, a comparison of the TF motif family organization in different eukaryotic species suggests an increase of redundancy of binding with organism complexity.

INTRODUCTION

Transcriptional regulation plays a crucial role in many biological processes, ranging from intracellular metabolic and physiologic homeostasis to cellular differentiation and development [1–3]. Alterations of regulatory specificity are among the major sources of phenotypic diversity and evolutionary adaptation [4–6].

Transcriptional regulation is mainly controlled by a class of DNA binding proteins known as transcription factors (TFs). Typically, a TF is able to recognize and bind to a specific set of similar sequences, centered around a consensus one. This set of sequences defines a binding motif, which can be synthetically represented in a Position Weight Matrices (PWM), fully describing the variability and the information content of the characteristic sequence pattern, under the assumption that each TF-DNA base interaction is independent [7, 8].

The combined regulatory actions of the set of expressed TFs on their target genes in a particular cell type define the so-called *regulatory network*. Deciphering how these regulatory networks and their components evolved has been a crucial quest over the last decades [9–16].

TFs are usually classified in families on the basis of the structural similarity of their DNA-binding domains (DBDs), i.e., the protein components that mediate the TF-DNA interaction [17]. However, this classification lacks the detailed information about the preferences of sequence binding, which is ultimately the feature that defines the set of potential target genes of a TF. Currently, thanks to the remarkable progress in the characterization and classification of the human repertoire of TFs and their PWMs [18], it is possible to introduce an organization of TFs into families based on their preferences of sequence binding (here called "motif families"), and to address quantitatively the evolutionary origins of this organization. This is the main goal of the present paper. In fact, we propose a simple model, based on the Birth-Death-Innovation (BDI) paradigm [19], to describe the evolution of TF binding preferences, and in particular to capture the mechanisms that shaped the TF distribution in motif families. The model will be used as a "neutral scenario" to identify the main evolutionary forces acting on the TF repertoire and on its regulatory strategies in complex eukaryotes.

To this end, it is also important to characterize a few relevant features of the eukaryotic repertoire of TFs, which will be important ingredients of our model. In discussing this issue, a comparison between eukaryotes and prokaryotes is fairly instructive.

The constant repertoire of DBDs in the post-metazoan era points to cis-innovation as the dominant source of PWM evolution.

In BDI models, the evolutionary mechanisms taken into consideration are birth (duplication) of a gene or a protein domain, death (deletion), and innovation (the acquisition of a new element). There are essentially two main ways in which an organism can acquire a TF with a new set of binding sequences, thus introducing in the dynamics of the PWMs (and thus of the motif families) a completely new class. It can arise from duplication of an existing TF followed by a mutation-induced divergence of its binding preferences (hereafter referred to as *cis-innovation*), or there can be a *de novo* creation or acquisition of a new DBD. However, this second possibility intuitively requires a much more complex scenario, involving the "birth" of a functional sequence from a non-coding background or an event of horizontal gene transfer, which seems to be extremely rare in higher eukaryotes [20, 21].

This observation can be substantiated by looking at the evolution of DBDs. DBDs have a different distribution among species in prokaryotes and eukaryotes. In-

*correspondence to: antonio.rosanova@gmail.com

dependent studies in bacteria [22], archaea [23], plants and animals [24–26] have consistently shown that DBDs are drawn from relatively small, anciently conserved repertoires. However, such conserved repertoires are superkingdom-specific, with prokaryotes and eukaryotes sharing only few DBDs [27, 28].

Prokaryotic conserved DBDs are present in all the species, with no clearly distinguishable partition within the major bacterial phyla [28], probably due to the extensive use of horizontal gene transfer. On the other hand, the conservation patterns of DBDs within the three main eukaryotic kingdoms (metazoa, fungi and plants) are much more well defined [28], thus allowing a rough reconstruction of their evolutionary history. A relative small portion of eukaryotic DBD families (29% according to [28]) are shared among all the three kingdoms, including homeobox families, zinc fingers, HLH and bZIPs.

Focusing on metazoa, there is evidence that most of the TFs classes originated at the dawn of the animal kingdom, with DBDs highly conserved over the last 600 million years [29–31]. Therefore, two major bursts of DBD family expansion can be identified: the first when eukaryotes branched off from prokaryotes [28] and the second at the common ancestral node between animals and fungi [30, 31]. In particular, the metazoa-specific wave of expansion and diversification of TFs is thought to have provided the basal toolkit for generating the metazoan multicellular condition through the process of development [17, 30, 31].

Following these observations, we shall assume the DBD repertoire of complex eukaryotes as essentially stable over the last 600 million years. Therefore, it is reasonable to assume a low rate of de novo innovation in the post-metazoan era, as confirmed by the very few appearances of new DBDs, while most of the innovation, and thus the regulatory network rewiring, seems to come from a duplication-divergence mechanism, or in other words from *cis-innovation*.

Evolution of the set of TFs versus evolution of their binding preferences.
It is well known that there is a precise relation between the total number of TFs of an organism and the genome size. In prokaryotes the number of TFs increases almost quadratically with the genome size [32, 33], while this increase is much slower in eukaryotes (power-law exponent $\simeq 1.2$) [34]. Several models have been proposed to explain this intriguing relation [35–38]. All these models essentially agree on the idea that the total number of TFs is related to the organism complexity, as roughly captured by its genome size. It can be argued that this value has been tuned to an optimal one to address in the most efficient way the regulatory needs of the organism. In fact, it has been observed that an upper bound must exist on the total number of TFs to ensure an optimal coding strategy in which misrecognition errors are minimized [39].

Since we aim to describe only the evolution of the TF regulatory strategies in complex eukaryotes, we shall assume that the mean number of TFs is essentially constant over time and stably close to the optimal value. In fact, the dynamics in which we are interested in is the evolution of the binding preferences of these TFs, which is presumably acting on a faster timescale with respect to the changes in the TF total number. This assumption of a separation of time scales is in line with the notion of punctuated equilibrium often implied in several evolutionary models [19]: long period of stasis are punctuated by short bursts of evolutionary activity that involve radical alterations of the duplication and elimination rates. Between these periods of drastic changes, the system seems to rapidly relax to equilibrium. The assumption of equilibrium naturally implies an approximate balance between the mechanisms generating an inflow and an outflow of genes. As a result, the dynamics of binding preferences will be described by a process of duplication, deletion and divergence, while the total number of TFs is essentially stable. Given the evidence that de novo innovation is basically negligible in higher eukaryotes, the equlibrium condition translates in equal deletion and duplication rates. In fact, cis-innovation can only change the binding preferences of already existing genes. On the contrary, BDI models usually studied in the context of protein domain evolution consider de novo innovation as the main source of innovation [19, 40].

This paper proposes a mathematical description belonging to the BDI class, but conveniently modified to model the evolution of TF PWMs in higher eukaryotes. We will use a specific classification of TFs in motif families, i.e. grouping TFs with similar binding preferences, as a tool to understand the TF organization in terms of regulatory strategies. As we shall see, the model describes remarkably well, *with only one free parameter*, the global evolution of motif families, suggesting that the abundance of most of them have been only marginally modulated by selective pressure. Indeed, the "core" of the distribution can be fairly well explained by a random stochastic process based on a scenario of neutral evolution. However, two macroscopic deviations from our random model can be identified, which can be associated to selective forces acting on the regulatory network. The detailed study of these deviations will allow us to shed some light on the complex evolution of the TF repertoire of higher eukaryotes. The organization of DBD families into more specific motif families will be also analyzed and compared with our null model predictions. Few classes of TF DBDs will be shown to differentiate in an "anomalous" way with respect to the apparently neutral average behavior, and these deviation find motivation in specific selective pressures already suggested to act on those TFs in independent studies. Finally, a comparison between the motif family structure in different species suggests a general trend of increasing "redundancy" in TF binding preferences with organism complexity.

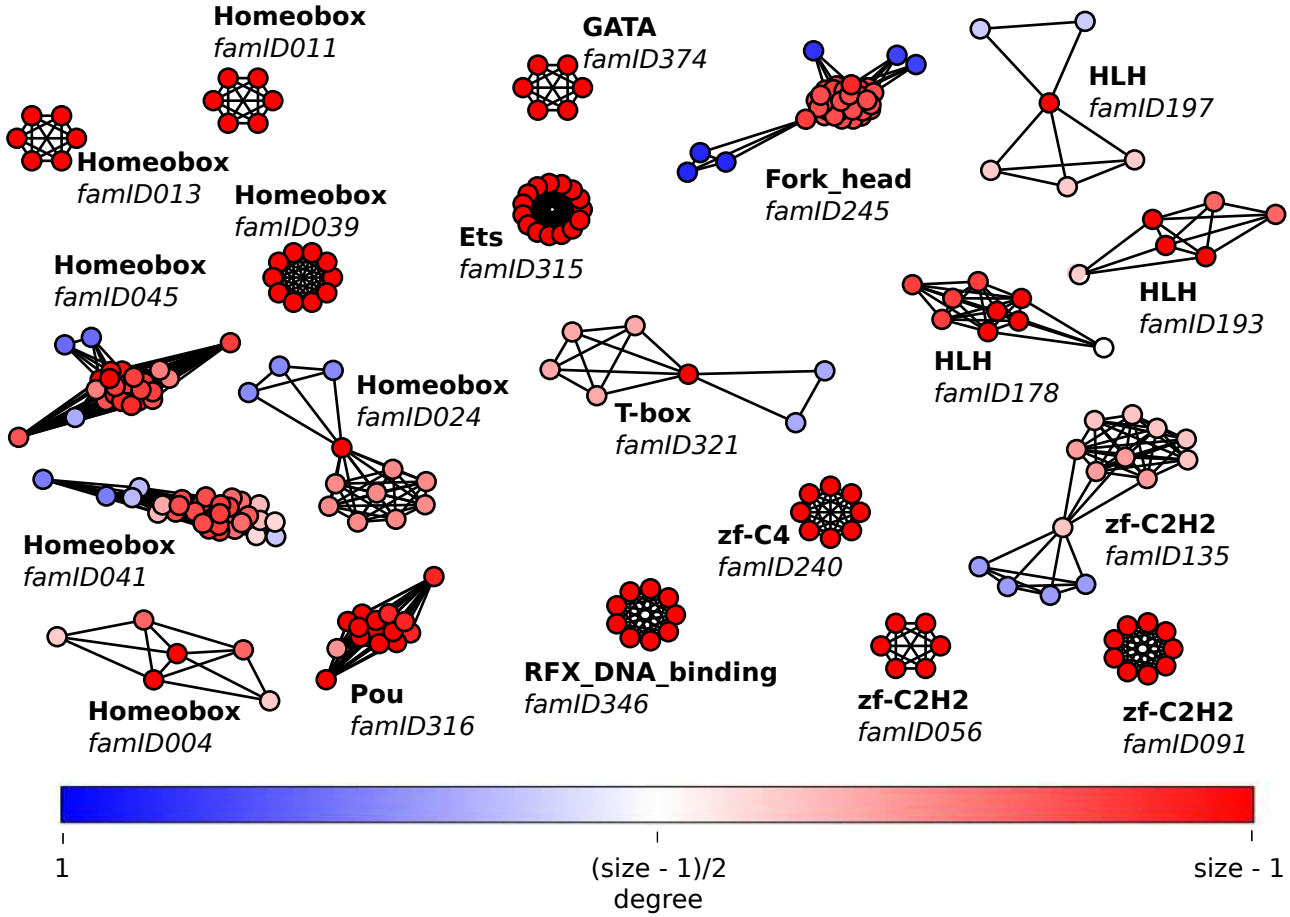

FIG. 1: **Graphical representation of motif families** with size > 5. Each vertex is a transcription factor, while links connect two TFs if they share at least one Position Weight Matrix ID. Colors identify the node degree with a color code (reported in the legend) spanning from blue, corresponding to degree 1, to red, corresponding to the maximal degree, i.e., the vertex is connected to all other nodes in the family. The circular layout highlights those families that are cliques. Each motif family is labeled with an internal ID code, see Supplementary Material for reference to family composition.

RESULTS

Organization of TFs in motif families

Large-scale studies [18, 41–44] have reported that TFs with similar DBD sequences tend to bind very similar DNA sequences, while TFs with different DBDs have clearly distinct specificities. More importantly, it has been shown that distinct TFs can be associated to a single PWM, which characterizes the binding preferences of the group [41]. Following [41], we propose a new classification of TFs based on PWMs, rather than on DBDs or on sequence homology. The result of this classification is an organization of TFs in what we call *motif families*, grouping together TFs with similar binding preferences (i.e., with the same PWM, see Materials and Methods). In order to obtain such motif families, we built a network where each gene is a vertex and an indirect edge connects two TFs if they are annotated with at least one common PWM. The network defined in this way is characterized by several disconnected components of high density, each of which defines a motif family. Figure 1 shows a snapshot of these families. Most of these components are cliques, i.e. groups of TFs with at least one PWM in common among all the members. Isolated TFs were considered as motif families of size 1. Figure 2 reports the size distribution of these families. It is interesting to compare this classification with the usual one based on DBDs. We verified that it never happens that different DBDs are contained in the same motif family, while most of the DBDs families are split in smaller more specific motif families. Indeed, the classification in terms of motif families represents a finer partition of the TF repertoire with respect to the usual one based on DBDs [41]. Moreover, the motif family classification is expected to be more closely related to the regulatory potential and thus to the functions of TFs. In fact, TFs with similar binding motifs show a high tendency to co-regulate the same target genes [39], and more generally

they show sets of target genes with similar functional annotations [18, 39, 41]. The empirical distribution of TFs in these motif families is supposed to be strictly related to the regulatory strategies of an organism, since changes in the PWMs associated to a TF imply alterations of its regulatory ability. The size distribution of these motif families in Figure 2 is the observable that we aim to explain in terms of a simple evolutionary model.

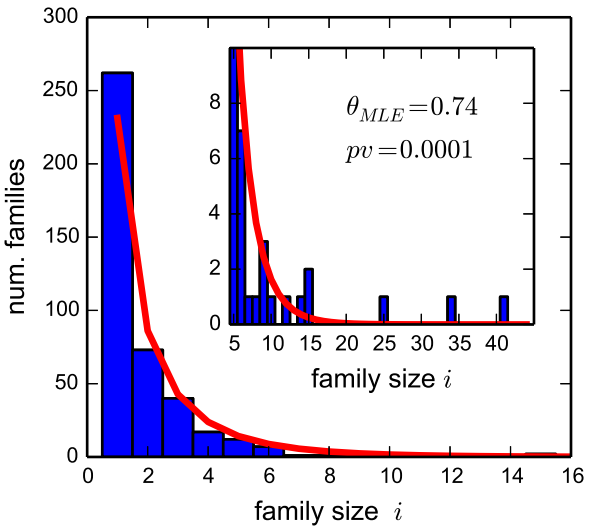

FIG. 2: **Size distribution of motif families in human.** The distribution accounts for 906 human TFs, organized in 424 families whose members share at least one PWM with at least another member. The inset is a zoom on the range of sizes > 5. The red-line is the best-fit model according to maximum likelihood estimation, which has a goodness-of-fit p-value $p < 0.0001$. The model captures the general trend, but clearly underestimates the number of families of size 1 and does not predict the presence of the largest families.

The Birth-Death-*cis*-Innovation model

The model we propose belongs to the general class of Birth-Death-Innovation models (for a thorough introduction see [40]). The focus of these models is on systems in which individual elements are grouped into families whose evolution is ruled by the dynamics of their individual members. These models typically include the elementary processes of family growth via element duplication (gene duplication), element deletion as a result of inactivation or loss (negative gene mutation), and innovation or emergence of a new family (neutral/positive gene mutation). All these processes are assumed to be of Markov type and the corresponding rates are assumed to be constant in time.

Applying this theoretical framework to the evolution of the TF repertoire and TF binding preferences in higher eukaryotes requires, as mentioned in the Introduction, a specific set of observation-based ingredients:

1. TFs evolve via gene duplication and divergence as in the typical BDI dynamics.

2. The repertoire of DBDs in higher eukaryotes is remarkably conserved over the last 600 million years. This implies that *cis-innovation* is the driving force of TF evolution on the time scale of PWM evolution we are interested in. In fact, our model description focuses only on the "late" stage of TF evolution in metazoans, in which very few new DBDs, and thus new motif families, are created de novo.

3. There is a separation between the time scale typical of the evolution of TF binding preferences and the time scale over which the number of TFs change significantly. This is essentially equivalent to the equilibrium hypothesis usually introduced in BDI models [19, 40], which implies that the family distribution is at a stationary state and the number of TFs is approximately constant.

Based on these ingredients, we propose a null-model of the BDI type to describe the specific features of the motif family distribution of TFs in higher eukaryotes (and in particular in the human case) and then use this model to identify deviations from this neutral scenario. To introduce the model in more detail, let us define as "class i" the set of all families of size i. Let f_i be the number of families in the i-th class, M be the total number of classes $i = 1....M$ (or the maximum size of a family), and N the total number of elements, thus representing also the extreme value for M. Acting at the "local" level on individual elements, the evolutionary dynamics shapes "globally" the system relocating a family from class i to class $i+1$ in case of duplication (or to class $i-1$ in case of removal).

Typically, BDI models [45–47] introduce innovation in the model only as a constant inflow in the class 1 due to de novo emergence of a new family (increase of f_1 by 1). As discussed above, we propose a generalization of the model by introducing also cis-innovation, in which an element of a family in class i mutates and gives rise to a new family. This results in the relocation of that element in class 1 and of its original family in class $i-1$ (i.e. a decrease of f_i and increase of f_{i-1} and f_1 by 1). Let λ, δ, ν and μ be the rates of element birth, death, *de novo*-innovation and *cis*-innovation respectively. Solving the master equations at the steady state (see the Materials and Methods section) one finds:

$$f_i = \frac{\nu + \mu N}{\lambda} \frac{\theta^i}{i}, \tag{1}$$

where $\theta = \frac{\lambda}{\delta+\mu}$.
The corresponding probability distribution p_i can be found straightforwardly by normalization:

$$p_i = \frac{f_i}{\sum_i f_i} = \frac{1}{\sum_i \frac{\theta^i}{i}} \frac{\theta^i}{i}. \quad (2)$$

A few comments are in order at this point:

- The normalized solution in Equation 2 gives a one-parameter prediction about the size distribution of motif families. The functional dependence on θ is equivalent to the one that can be obtained with standard BDI models [40], i.e., with de novo innovation as the only source of innovation. However, our generalized model suggests a different interpretation of the parameter. In fact, $\theta = \frac{\lambda}{\delta+\mu}$ and thus its value depends on the rate of cis-innovation.

- The steady state condition is $\frac{df_i}{dt} = 0 \; \forall i$, or equivalently the total number of elements $N = \sum_i^M i f_i$ is constant over time. This condition translates into the parameter constraint $\lambda + \nu = \delta$.

- As previously discussed, we expect ν to be very small in our case (i.e., negligible de novo innovation), and accordingly we shall approximate $\nu \to 0$ in the following. We shall further verify "a posteriori" the validity of this approximation using an independent analysis on the evolution of TFs in different lineages (see below). In this regime, the stationary condition simplifies to a balance between duplication and deletion rates $\lambda = \delta$, and $\theta \simeq \frac{1}{1+\mu/\lambda}$. Therefore, the deviation of θ from 1 allows to directly estimate the magnitude of μ with respect to λ, i.e., the relevance of cis-innovation with respect to the birth/death rate. As we will see below, a comparison with the data in the human case supports a value of $\theta \sim 0.73$, thus highlighting the important role that cis-innovation had in the recent evolution of the eukayotic TF repertoire. Moreover, within this approximation, also the family distribution in Equation 1 can be written in a very simple and compact form:

$$f_i = N\frac{\mu}{\lambda}\left(\frac{\lambda}{\delta+\mu}\right)^i \frac{1}{i} = N(1-\theta)\frac{\theta^{i-1}}{i}. \quad (3)$$

- An analytical estimate of the number of classes $F = \sum_i^M f_i$ in which the N elements are organized when the dynamics reaches equilibrium can also be calculated as:

$$\frac{F}{N} = \frac{1-\theta}{\theta}\sum_i^M \frac{\theta^i}{i} \simeq \frac{\theta-1}{\theta} ln(1-\theta) \quad (4)$$

This represents the neutral model prediction on the number of motif families given a set of N TFs subjected to the described BDI dynamics.

The model can explain the core of the size distribution of motif families and identifies two main deviations

We compared the distribution predicted by our model, Equation 2, with the empirical TF organization in motif families obtained from the CIS-BP database (Materials and Methods). This comparison has been quantified by estimating a best fit value for θ with a Maximum Likelihood method and a p-value associated to the quality of the fit using a *goodness-of-fit* test based on the Kolmogorov-Smirnov statistics (Materials and Methods). Although the central part of the size distribution seems well captured by the theoretical model, a direct fit of the whole distribution gives very low p-values (p-value $< 10^{-3}$, see Figure 2). This poor p-value shows the presence of significative deviations with respect to our random null-model. These deviations can be easily identified looking at Figure 2. They are located at the two ends of the distribution and involve a few of the largest families and the smallest ones (i.e., families of size 1). Using the KS test and a p-value threshold for acceptance of 0.75, we can identify in a quantitatively and consistent way the fraction (about 25%) of isolated TFs and the number (three) of the largest families which account for most of the deviations from the null model (Materials and Methods and Figure 3).

If we subtract from the whole distribution these two tails (for a total of $\sim$ 150 TFs, i.e. about 16% of the total number of TFs in analysis), we eventually find a remarkable agreement between the model predictions and experimental data (p-value ~ 0.8, see Figure 4). Therefore, the "core" of the distribution is well described by the exponential-like solution of Eq. (3), while deviations are due to few families that can be isolated and studied in detail. This suggests that the evolution of a large portion of the TF repertoire in higher eukaryotes was driven by a neutral stochastic process of the BDI type with only two exceptions: an excess of isolated TFs and three large families which on the contrary are characterized by a strong level of duplication without innovation. Let us address in more detail these two deviations.

Isolated transcription factors

The fitting procedure allows us to obtain a rough estimate of the fraction N_s of size 1 families which are not explained by our theoretical description. This number is in the range $40 < N_s < 80$, i.e, in between 20% and 30% of the total number of size 1 families (Materials and Methods and Figure 3). The emergence of a size 1 family in our model description can come from de novo innovation or from duplication of an existing TF, followed by a cis-innovation event that defines a new PWM. We argued that de novo innovation is negligible in our case of study, so we expect that most of the isolated TFs are the result of a previous duplication event. In this scenario,

they should share their DBD at least with the TF they duplicated from, and we verified that indeed empirically this is the case for the majority of isolated TFs, thus supporting our model description. However, some isolated TFs have a DBD which is not shared with any other TF (12 in our sample) or are characterized by a DBD which is classified as 'UNKNOWN' (44 in our sample), so also potentially unique. The presence of these isolated TFs with unique DBDs can be explained by the two following mechanisms.

Newly acquired DBDs. A few of them are due to actual recent de novo innovation events, thus introducing new DBDs in the last period of post-metazoan evolution. These "recent" TFs appear in our analysis most likely as size 1 families only because they had not time to enter into the duplication process. Looking at the orthology maps we can rather easily identify these DBDs and the corresponding TFs (see Supplementary Material and below) which turn out to be very few, thus supporting "a posteriori" our $\nu = 0$ approximation.

Singleton genes. The majority of excess isolated TFs are most probably *singleton genes* for which duplication is peculiarly avoided. The existence of this class of genes has been recently proposed [48, 49]. They are supposed to be ancestral genes of prokaryotic origin, addressing basilar functions and requiring a fine-tuning of their abundances, thus making their duplication particularly detrimental. They would be the result of a selective pressure to avoid duplication, and thus, by definition, cannot be explained by our neutral model.

Since singleton genes are not included in our model, they are good candidates to explain the excess of isolated TFs in Figure 2. To distinguish between putative singleton genes and recent genes in the motif families of size 1, we analyzed their evolutionary origin. More specifically, we manually inspected the taxonomic profiles of these 56 TFs in the EggNOG database [50]: 16 of them have a putative origin at the Last Universal Common Ancestor (LUCA), i.e. they are shared among bacteria, archaea and eukarya; 25 are in common among all eukarya, 4 among opisthokonta, 3 among metazoa and 8 have a post-metazoan origin. Therefore, at least 41 of these TFs have a very ancient origin (LUCA + eukarya) and could well be examples of "singleton" TFs, while 8 are instead of very recent origin (post-metazoan, but 4 of them are shared only among euteleostomi) and are thus likely to be "recent" TFs. These recent TFs constitute

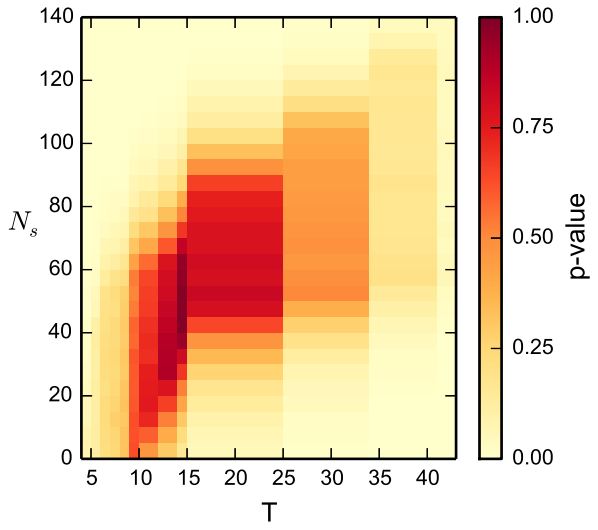

FIG. 3: **Heatmap for the goodness-of-fit p-value as the data sample is reduced**. On the x-axis, T indicates the threshold in size above which families are excluded from the sample. On y-axis N_s indicates the number of families of size one excluded from the sample. An increase in T or N_s reduces the sample size in analysis by reducing the number of TFs considered. For each sample size a goodness-of-fit test for the best-fit model was performed and the corresponding p-value is reported with the color code in the legend. Considering a p-value of 0.75 as the acceptance limit identifies $T = 25$ as the size threshold at which the fit is acceptable. This corresponds to the exclusion of the three largest families. For such a threshold T, the optimal values for the p-value are reached for values of N_s in the range $40 < N_s < 80$.

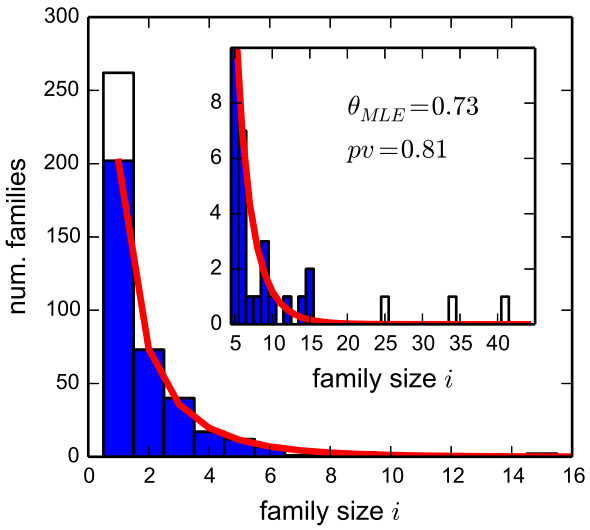

FIG. 4: **Size distribution of motif families for the reduced sample**. The filled distribution represents the motif family size distribution for the reduced dataset, while the original empirical distribution of Figure 2 is reported with the unfilled bars. The inset shows a zoom on the range of sizes > 5. The line represents the prediction of our model with the best fit choice of the parameter θ , which turns out to fit very well the data contained in the reduced sample with a goodness of fit p-value $p \sim 0.8$. The best fit value $\theta = 0.73$ does not differ substantially from the value that is obtained by fitting the whole empirical sample as in Figure 2.

less than the 1% of our sample, supporting "a posteriori" the $\nu = 0$ approximation.

To find additional evidence that these 41 ancient TFs can be bona fide "singleton genes", we queried the NGC5.0 database [51], which provides information about the gene duplicability for a large set of cancer genes. 14 of our putative singletons are present in this collection, and 12 of them show indeed no evidence of duplicability (at 60% coverage), thus supporting their "singleton" nature. It is interesting to notice that the overall number of putative singletons (41 genes) is compatible to the size of the deviation from the random null model ($40 < N_s < 80$) observed in our best fit tests.

Over-expanded families

Our analysis singles out also three over-expanded families. The over-expansion can be due to two parallel mechanisms: an enhanced rate of duplication and/or a decreased rate of cis-innovation. Looking at the three over-expanded families, three very homogeneous groups of TFs can be recognized: the FOX family (size 41), the HOX family (size 34) and another homeobox family (size 25). These three families are good examples of the two mechanisms mentioned above. The HOX family contains TFs well known for their role in morphogenesis and animal body development [52]. Also TFs in the other over-expanded homeobox family show enrichments for GO annotations related to *morphogenesis*, *development* and *pattern specification*, as reported in Table I. These two families may well represent cases of positive selection for duplication and subsequent fixation. Due to their crucial role in morphogenesis, these TFs could have been under a strong conservation pressure that inhibited mutations leading to gene loss after duplication. The third family, which is the largests one, collects most of the FOX (Forkhead box) TFs present in the sample. TFs belonging to this family are known to be "bispecific", i.e. they recognize two distinct DNA sequences [53], and for this reason they play an important and peculiar role in the regulatory network of metazoans [53]. While their over-expansion can be due to positive selection for functional reasons, their unique feature of bispecific binding could suggest that innovation is particularly difficult for these TFs. In fact, bispecifity is likely to impose stronger constraints, from a structural point of view, than those imposed on other TFs. In this perspective, it is interesting to stress the different distribution of Forkhead and Homeobox genes in motif families. Almost all the Forkhead genes are collected in this single large motif family, suggesting no cis-innovation events that would have moved some of these genes in families of other sizes. Only 6 Forkhead TFs are present in other motif families. On the other hand, Homeobox genes, besides the two main families discussed above, are dispersed in several other motif families, thus are associated to a variety of PWMs. This difference suggests that duplication of Homeobox genes has been positively selected at a certain time point probably because of their crucial role in the development of multicellular organisms (see Table I), but cis-innovation have progressively changed their binding preferences. On the other hand, very few events of cis-innovation are associated to FOX genes that indeed "accumulated" in a single motif family. These interpretations of the possible evolutionary origins of the over-expanded motif families will be addressed in more detail in the next section.

Phenomenology of the splitting of DBD families in motif families

So far, we considered the "global" distribution of all TFs in motif families. However, the evolution of different DBD families can be considered as independent, and thus can be studied separately. In other words, we can focus on the dynamics of the splitting of each DBD family in smaller motif families through the process of cis-innovation followed by duplication and deletion. Indeed, the dynamics of duplication and deletion can only expand or reduce a certain DBD family size, while cis-innovation can change the binding preferences of a TF, thus creating new motif families, but it is not expected to create a brand new DBD class. We provided evidence that the DBD repertoire has remained essentially constant on the evolutionary time scales here in analysis. Therefore, our model can also be used to predict how, in a neutral scenario, each DBD family is expected to split in different motif families through the cis-innovation process. Our model directly provides an analytical estimate of the average number of families in which a set of TFs is expected to be organized in Equation (4). To evaluate also the possible variability of this neutral expectation, we ran $5*10^4$ simulations of the model for different system sizes, corresponding to the different numbers of TFs in the DBD families. The rates of duplication, deletion and cis-innovation are set by the fit of the global distribution of motif families in Figure 2 ($\theta = 0.74$) and by the additional constraint $\lambda = \delta$. In a neutral scenario, we expect the rates to be the same for all DBD families. Clearly, considering the DBD families separately significantly reduces the statistics that can be used, but we can take advantage of a parameter-free prediction of the model (fitted on the whole dataset) to compare with. Figure 5 shows the numbers of motif families in which a DBD is divided, given its number of TFs. The analytical prediction (dashed line) and the results of model simulations (shaded areas represent 1 and 3 standard deviations from the average simulated behaviour) are compared with empirical data (symbols). While most of the splitting patterns of DBDs families do not deviate significantly from the model prediction, three clear "outliers" can be observed. The Forkhead DBD family shows a smaller than

GO biological process complete	Fold Enrichment	p-value
embryonic skeletal system development (GO:0048706)	6.92	7.87E-18
skeletal system development (GO:0001501)	4.21	8.44E-15
embryonic skeletal system morphogenesis (GO:0048704)	7.18	1.20E-14
skeletal system morphogenesis (GO:0048705)	5.77	4.29E-14
anterior/posterior pattern specification (GO:0009952)	4.91	2.55E-13
pattern specification process (GO:0007389)	3.10	3.15E-10
regionalization (GO:0003002)	3.35	7.25E-10
embryonic organ morphogenesis (GO:0048562)	3.47	4.71E-08
embryonic morphogenesis (GO:0048598)	2.66	1.83E-06
organ morphogenesis (GO:0009887)	2.25	4.76E-06
chordate embryonic development (GO:0043009)	2.59	7.91E-06
embryo development ending in birth or egg hatching (GO:0009792)	2.58	9.09E-06
embryo development (GO:0009790)	2.10	2.08E-05
embryonic organ development (GO:0048568)	2.65	4.30E-05
anatomical structure morphogenesis (GO:0009653)	1.76	5.73E-05
system development (GO:0048731)	1.43	7.57E-04
anatomical structure development (GO:0048856)	1.33	8.57E-04
multicellular organism development (GO:0007275)	1.36	1.19E-03
animal organ development (GO:0048513)	1.49	1.97E-03
developmental process (GO:0032502)	1.30	2.77E-03
single-multicellular organism process (GO:0044707)	1.32	4.44E-03
multicellular organismal process (GO:0032501)	1.30	1.25E-02
cranial nerve morphogenesis (GO:0021602)	7.94	1.36E-02
single-organism developmental process (GO:0044767)	1.29	1.58E-02
cranial nerve development (GO:0021545)	5.31	3.00E-02
nerve development (GO:0021675)	5.13	4.04E-02

TABLE I: **Gene Ontology analysis** of the genes belonging to the two homeobox families of size 25 and 34. Only results with p-value below 0.01 are shown. The enriched annotations are mainly associated to development and morphogenesis of multicellular organisms.

expected number of motif families. We have previously show that the largest motif family including most of the Forkhead associated TFs is also over-expanded. More specifically, most of the Forkhead TFs belongs to one single motif family. We speculated that the overexpansion could have been driven by structural constraints limiting the evolvability of the binding preferences. Indeed, the small number of motif families in which this DBD class divides into can be indicative of a lower-than-average rate of cis-innovation. Similarly, the Homeobox DBD family appear to be partitioned in less motif families than expected, and as previously discussed, only some of these families appear to be over-expanded, probably for specific functional reasons.

The other significant deviation from the model prediction concerns a Zinc Finger class of TFs, that appears to have greatly diversified the TF PWMs. The corresponding motif families are not over-expanded, in fact they did not emerge as deviations in the previous analysis (Figure 2). In fact, the histogram of their motif family sizes (the analogous of Figure 2 but restricted to this Zinc Finger TFs) follows reasonably well our null model. However, the fitted parameter $\theta = 0.56$ is well below the value obtained from the all set of TFs ($\theta = 0.74$), thus confirming again that the rate of cis-innovation for this DBD family is higher than the average rate for all TFs.

Zinc Finger TFs are known to be characterized by multiple tandem C2H2 zinc finger domains. Such modularity enabled a rapid functional divergence among recently duplicated paralogs, as each domain in the protein can be gain, loss and mutate independently [54]. This piece of evidence seems to further support a higher-than-average tendency of TFs in this DBD class to diversify their binding preferences.

TF redundancy of binding increases with organism complexity

This section addresses the differences in the motif family organization in different eukaryotic species. In particular, we focused on model species, which are expected to have well annotated TF repertoires. The same type of analysis presented in Figure 2 was performed on the set of TFs of yeast and of three other species of increasing complexity in the animal lineage: *C. elegans*, *D. melanogaster* and *M. musculus*. Figure 6 shows the histograms of the family size distributions and the corresponding fits with the null model in Equation (2). In all tested cases, the motif families distribution follows the model functional form with an agreement comparable to the human case. However, there is a trend in the fitted parameter θ with complexity as measured by the number of TFs in the species (or alternatively by the total number of genes).

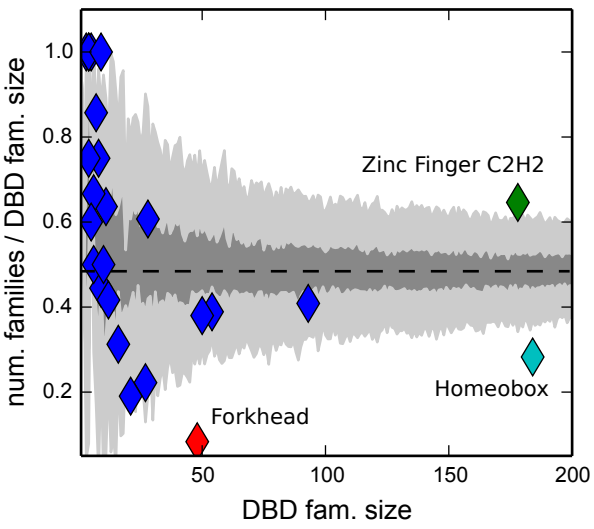

FIG. 5: **Splitting of DBD families in motif families.** The numbers of motif families in which a DBD class is divided into is shown as a function of the number of TFs in that class. Each point represents the empirical value for a DBD family, while the dashed line represents the expected value for $\theta = 0.74$ as given by Eq. 4. In order to evaluate the fluctuations on the expectation, we simulated the evolution of $5 * 10^4$ DBD families, with starting size ranging from 1 to 500, $\theta = 0.74$ and $\lambda = \delta$. The two shaded areas correspond to 1 standard deviation and 3 standard deviations from the average. Green diamond: Zinc Finger C2H2 family. Cyan diamond: Homeobox family. Red diamond: Forkhead family.

This trend is reported in Figure 6 and it is sublinear in the investigated window of TF repertoires. The definition of $\theta \simeq \frac{1}{1+\mu/\lambda}$ indicates that this trend corresponds to a decrease rate of cis-innovation, with respect to the duplication rate, as the complexity of the organism increases. The value of θ intuitively represents the level of "redundancy", i.e., the tendency of TFs to keep the same binding preferences. Indeed, for the limit value $\theta = 1$ the distribution in Equation (1) becomes a power-law distribution, thus with an increasing number of motif familes with a large number of TFs. Figure 6 shows that this level of "redundancy" increases with the organims complexity. As we will argue in the Discussion, this result can be linked to the known increased degeneracy of the TF binding sites in higher eukaryotes [27].

DISCUSSION

A comparison of metazoan lineages shows that the ancestral metazoan genome included members of the bHLH, Mef2, Fox, Sox, T-box, ETS, nuclear receptor, bZIP and SMAD families, and a diversity of homeobox-containing classes, including ANTP, Prd-like, Pax, POU, LIM-HD, Six and TALE [29, 30, 55–58]. This implies that most of the human DBDs originated at the dawn of the animal kingdom, before the divergence of contemporary animal lineages, providing a genetic toolkit conserved from the bilateria lineage on. Once the main DBDs were established, the discovery of new ones dropped down and the lineage speciation took advantage of a finer divergence dynamics at the regulatory level, by changing just the DNA-sequence preferences of binding within the DBD.

Following this observation, this paper proposes a new way to organize transcription factors into families, which we call *motif families*, based on their PWMs. It also introduces a simple one-parameter model of the BDI type to explain the distribution of TFs in these motif families. The novelty of this model with respect to standard BDI models is the introduction of what we called $cis-innovation$. Cis-innovation accounts for a mutation-induced change of the binding preferences of an existing TF, thus implying a rewiring in the regulatory network but without the introduction of a new TF associated to a new PWM, which seems an extremely rare event on the time scales in analysis. Using this simple theoretical description, we showed that the recent evolution of the majority (more than 80%) of the human TFs is rather well described by a neutral stochastic process of duplication, deletion and cis-innovation. Furthermore we devised two main deviations from this neutral scenario: the overexpansion of three motif families and the "freezing" of a set of isolated TFs (i.e., size one motif families).

The deviations which we observed, involving more or less 20% of the TFs, seem to be due to opposite evolutionary pressures. The three over-expanded families are associated to an enhanced duplication rate and/or a decreased cis-innovation rate. The excess of isolated TFs seems instead to be mainly due to the inhibition of duplication for a specific set of ancient TFs, or "singletons". The largest of the over-expanded families, collects most of the FOX TFs (Forkhead box) present in the TFs in analysis. The second and third over-expanded families contain TFs belonging to the Homeobox family (one of which contains most of the HOX TFs). These last two families seems to represent examples of positive selection for duplication due to the crucial role their TFs play in morphogenesis and animal body development [52]. Indeed, a gene ontology analysis for their genes with respect to the entire set of TFs (see Materials and Methods and Table:I) shows enrichments for GO annotations related to *morphogenesis*, *development* and *pattern specification*. This result strongly supports the idea that the expansion of these families was positively selected for their crucial role in multicellularity evolution [30] and animal speciation. [52]. The overexpansion of FOX genes can also be associated to positive selection because of their importance in the development of several metazoans [53], but probably coupled with a functional constraint related to the bispecificity of their DNA binding limiting the rate of allowed cis-innovation. This low rate of cis-innovation was further confirmed by the low level of diversification of the Forkehad DBD family into motif families presented

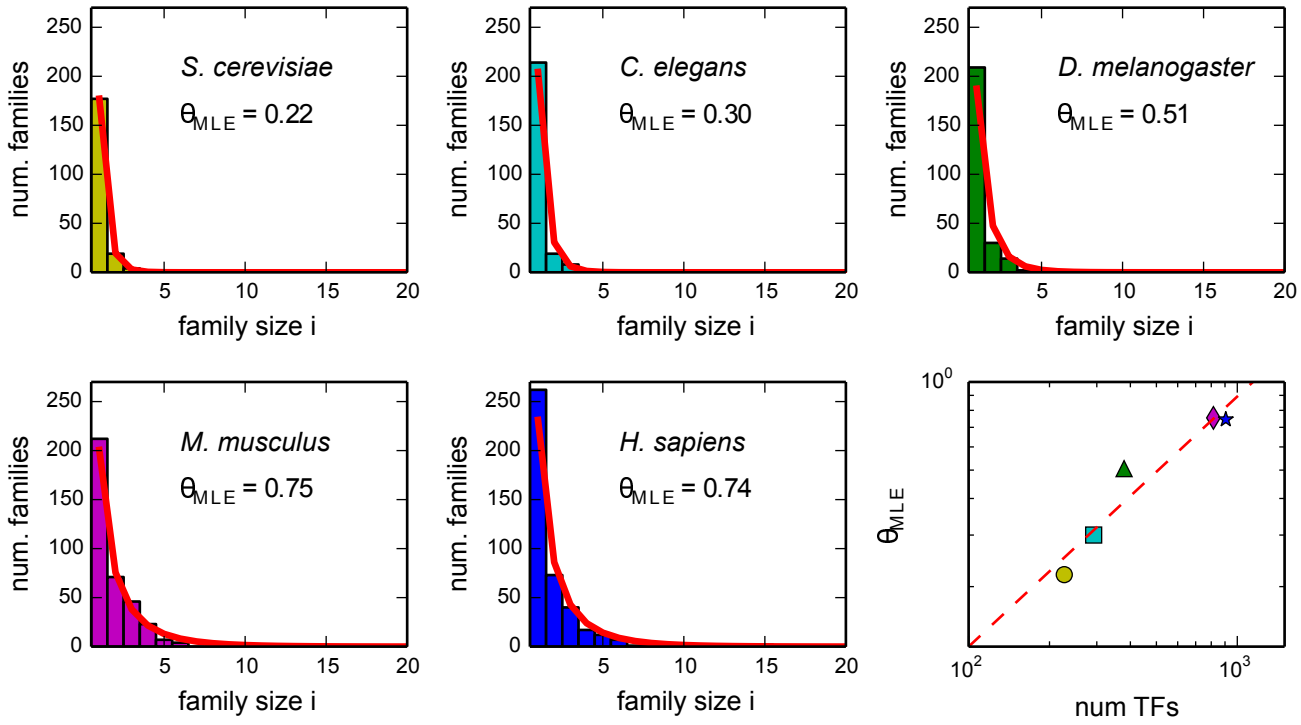

FIG. 6: **Size distribution of TF motif families for different eukaryotic organisms.** We report the distributions and the best fit values of θ for five different organisms of increasing complexity. As for the human case, the data are taken from the CIS-BP database. The bottom-right panel shows how θ scales with the number of TFs. The red line is the fit $\theta \sim (\#\ \text{TFs})^{0.85}$.

in Figure 5. The analysis of the splitting of the DBD families into motif families also revealed a higher-than-average rate of diversification of the Zinc Finger DBD family. A result that seems compatible with the high level of "evolvability" of these TFs suggested in an independent study [54].

We also detected an over-presence of families of size 1, identifying a portion of $40 < N_s < 80$ TFs as deviation from the null model. We found a comparable number of TFs (41) that have a very ancient origin and no evidence of duplication, and thus are likely to be examples of the so-called *singleton genes* [48, 49]. According to the "singleton hypothesis", these genes are indeed ancestral genes of prokaryotic origin, addressing basic functions and requiring a precise fine-tuning of their abundances [49]. Therefore, the effect of selection against duplication seems to be perfectly compatible with their proposed functional role.

A major issue in the study of the evolution of regulatory networks is to identify those features of the network which can be in some way associated to the organism complexity. Combinatorial regulation is a distinctive feature of complex eukaryotes. Indeed, prokaryotic and eukaryotic TFs use different binding strategies, with PWMs of high and low information content respectively [27]. This difference is related to the evolution of the combinatorial strategies of control, typical of higher eukaryotes, that can compensate the low information content of their TF binding sites [27] by having specific combinations of TFs targeting the same promoter. This could have also been favoured by the widespread presence of transposable elements able to convey combinations of TF binding sites all over the genome [59].

Accordingly, we expect a tendency of eukaryotes of increasing complexity to increase their TF repertoire, as it can indeed be observed [34], to have a richer repertoire to implement combinatorial regulation. However, the increased degeneracy of TF PWMs with the organism complexity can also have another relevant consequence. If the set of preferred binding sequences defining the motif family of a TF is loosely defined, it can include several possible sequences. Thus, the mutation process is less likely to drive the TF away from its motif family. This would translate in a lower cis-innovation rate in our model for organisms with higher complexity, and this trend seems indeed to emerge from our comparison of the different motif family organization in different species (Figure 6). Complexity seems to be associate to the "redundancy" of the TF repertoire, i.e., to the presence of large families of TFs which recognize the same binding sequences. It would be interesting to understand the consequences of this observation on the topology and function of the

regulatory network.

MATERIALS AND METHODS

Data set

We took advantage of the Catalog of Inferred Sequence Binding Preferences (CIS-BP database) [18], which collects the specificities of a vast amount of TFs in several species. The PWMs in this database were either directly derived from systematic protein binding microarray (PBM) experiments or inferred by overall DBD amino acid identity. Furthermore, the CIS-BP database gathers data from all the main existing databases (such as TRANSFAC [60], JASPAR [61] and SELEX [62]) and several Chip-Seq experiments, which had been used for cross-validation. To construct the motif families, we downloaded the PWMs associated to each TF, considering both those obtained from experimental assays and the inferred ones. In this way, we obtained 4172 PWM unique identifiers (PWD IDs) annotated to 906 different TFs.

The BDI model

We define as "class i" the set of all families of size i. f_i represents the number of families in the i-th class and M be the total number of classes $i = 1....M$ corresponding to the possible family sizes, with M at most equal to the total number of elements N.

The evolution equations are:

$$\begin{aligned}\frac{\mathrm{d}f_1(t)}{\mathrm{d}t} =& -(\lambda+\delta+\mu)f_1(t) + 2(\delta+\mu)f_2(t) + \mu N + \nu \\ \frac{\mathrm{d}f_i(t)}{\mathrm{d}t} =& (i-1)\lambda f_{i-1}(t) - i(\lambda+\delta+\mu)f_i + \\ &+ (i+1)(\delta+\mu)f_{i+1}(t) \\ \frac{\mathrm{d}f_M(t)}{\mathrm{d}t} =& (M-1)\lambda f_{M-1}(t) - M(\delta+\mu)f_M(t)\end{aligned} \tag{5}$$

where λ, δ, ν and μ denote the birth, death, *de novo* innovation and *cis*-innovation rates respectively.

The model can be mapped in the simplest case of the BDI models discussed in [40] with the substitution $\delta' = \delta + \mu$ and $\nu' = \nu + \mu N$.

From the general solution discussed in [40], we obtain at steady state:

$$f_i = \frac{\nu'}{\lambda}\left(\frac{\lambda}{\delta'}\right)^i \frac{1}{i} \sim \frac{\theta^i}{i} \tag{6}$$

where $\theta = \frac{\lambda}{\delta'} = \frac{\lambda}{\delta+\mu}$. If, following [40], we assume a balance between birth and death rates $\lambda = \delta$ then $\theta = \frac{\lambda}{\lambda+\mu}$ and Eq.(6) becomes:

$$f_i = \frac{\nu+\mu N}{\lambda}\left(\frac{\lambda}{\lambda+\mu}\right)^i \frac{1}{i} \tag{7}$$

The deviation of θ from 1 allows to estimate the magnitude of μ with respect to λ. In the limit of $\theta \to 1$ ($\mu \to 0$) the usual power-like behaviour of the standard DBI model is recovered. Since we know $\nu \ll \mu$, we shall assume $\nu = 0$ and the solution of the model eq.(7) becomes a function only of θ.

$$f_i = N\frac{\mu}{\lambda}\left(\frac{\lambda}{\lambda+\mu}\right)^i \frac{1}{i} = N(1-\theta)\frac{\theta^{i-i}}{i}$$

Maximum Likelihood estimation of θ

To perform a MLE of the parameter θ, we must first move from the distribution of the number of families to a probability distribution. This is simply achieved by normalizing the f_i. $p_i = C_M\frac{\theta^i}{i}$. The normalization constant C_M assumes a very simple form in the large M limit:

$$C_M = [\sum_{i=1}^{M}\frac{\theta^i}{i}]^{-1} \overset{M\to\infty}{=} [-\ln(1-\theta)]^{-1},$$

leading to the probability distribution:

$$p_i = \frac{1}{-ln(1-\theta)}\frac{\theta^i}{i} \tag{8}$$

We show in the Supplementary Material that for our range of values of M and θ the error induced by this approximation is negligible.

The probability distribution in Eq. 8 is simple enough to allow an analytic determination of the MLE for θ (see the Supplementary Material for the detailed calculation), which turns out to be:

$$\theta_{MLE} = 1 - e^{\frac{1}{k}+W_{-1}\left(-\frac{1}{k}e^{-\frac{1}{k}}\right)} \tag{9}$$

Where $\overline{k}$ is the mean size over the sample and W is the Lambert Function.

Goodness-of-fit test

We compared the empirical data with our model, defined by θ_{MLE}, following the strategy proposed in [63]. More precisely we used the Kolmogorov-Smirnov (KS) statistic as a measure of the distance between the distribution of the empirical data and our model. In order to obtain an unbiased estimate for the p-value, we created a set of one thousand synthetic data samples with the same size of the empirical one, drawn from a distribution with

the same θ_{MLE} value. For each synthetic sample, we computed the KS statistic relative to the best-fit law for that set and constructed the distribution of KS values. The p-values reported in the paper represent the fraction of the synthetic distances larger than the empirical one.

Gene Ontology

We performed a gene ontology analysis on the genes belonging to the union of the three larger motif families using the overrepresentation test of the PANTHER facility [64] and selecting only the Biological Process ontology. We chose as a background for the test the entire data sample (906 TFs) to eliminate annotations simply associated to generic regulatory functions of TFs. p-values were evaluated using the Bonferroni correction.

FUNDING

The work was partially supported by the Compagnia San Paolo grant GeneRNet.

ACKNOWLEDGEMENTS

We thank F.D. Ciccarelli and M. Cosentino Lagomarsino for critical reading of the manuscript, and A. Colliva, M. Fumagalli and A. Mazzolini for useful discussions. The authors declare no competing financial interests.

[1] Domenico Accili and Karen C Arden. Foxos at the crossroads of cellular metabolism, differentiation, and transformation. *Cell*, 117(4):421–426, 2004.

[2] Gretchen Bain, Els C Robanus Maandag, David J Izon, Derk Amsen, Ada M Kruisbeek, Bennett C Weintraub, Ian Krop, Mark S Schlissel, Ann J Feeney, Marian van Roon, et al. E2a proteins are required for proper b cell development and initiation of immunoglobulin gene rearrangements. *Cell*, 79(5):885–892, 1994.

[3] Brian D Dynlacht. Regulation of transcription by proteins that control the cell cycle. *Nature*, 389(6647):149–152, 1997.

[4] Carlos D Bustamante, Adi Fledel-Alon, Scott Williamson, Rasmus Nielsen, Melissa Todd Hubisz, Stephen Glanowski, David M Tanenbaum, Thomas J White, John J Sninsky, Ryan D Hernandez, et al. Natural selection on protein-coding genes in the human genome. *Nature*, 437(7062):1153–1157, 2005.

[5] Subhajyoti De, Nuria Lopez-Bigas, and Sarah A Teichmann. Patterns of evolutionary constraints on genes in humans. *BMC Evol. Biol.*, 8(1):275, 2008.

[6] Nuria Lopez-Bigas, Subhajyoti De, and Sarah A Teichmann. Functional protein divergence in the evolution of homo sapiens. *Genome Biol*, 9(2):R33, 2008.

[7] Panayiotis V Benos, Martha L Bulyk, and Gary D Stormo. Additivity in protein–dna interactions: how good an approximation is it? *Nucleic Acids Res.*, 30(20):4442–4451, 2002.

[8] Emmanuelle Roulet, Stéphane Busso, Anamaria A Camargo, Andrew JG Simpson, Nicolas Mermod, and Philipp Bucher. High-throughput selex–sage method for quantitative modeling of transcription-factor binding sites. *Nature biotechnology*, 20(8):831–835, 2002.

[9] Sarah A Teichmann and M Madan Babu. Gene regulatory network growth by duplication. *Nature genetics*, 36(5):492–496, 2004.

[10] M Madan Babu, Sarah A Teichmann, and L Aravind. Evolutionary dynamics of prokaryotic transcriptional regulatory networks. *J. Mol. Biol.*, 358(2):614–33, apr 2006.

[11] Otto X Cordero and Paulien Hogeweg. Feed-forward loop circuits as a side effect of genome evolution. *Molecular biology and evolution*, 23(10):1931–6, oct 2006.

[12] Jakob Enemark and Kim Sneppen. Gene duplication models for directed networks with limits on growth. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(11):P11007–P11007, nov 2007.

[13] John W Pinney, Grigoris D Amoutzias, Magnus Rattray, and David L Robertson. Reconstruction of ancestral protein interaction networks for the bzip transcription factors. *Proc. Natl. Acad. Sci. U. S. A.*, 104(51):20449–53, dec 2007.

[14] Maximino Aldana, Enrique Balleza, Stuart Kauffman, and Osbaldo Resendiz. Robustness and evolvability in genetic regulatory networks. *Journal of theoretical biology*, 245(3):433–48, apr 2007.

[15] Anton Crombach and Paulien Hogeweg. Evolution of evolvability in gene regulatory networks. *PLoS Comput. Biol.*, 4(7):e1000112, 2008.

[16] Katja Nowick and Lisa Stubbs. Lineage-specific transcription factors and the evolution of gene regulatory networks. *Briefings in functional genomics*, 9(1):65–78, jan 2010.

[17] Juan M Vaquerizas, Sarah K Kummerfeld, Sarah A Teichmann, and Nicholas M Luscombe. A census of human transcription factors: function, expression and evolution. *Nat. Rev. Genet.*, 10(4):252–263, 2009.

[18] Matthew T Weirauch, Ally Yang, Mihai Albu, Atina G Cote, Alejandro Montenegro-Montero, Philipp Drewe, Hamed S Najafabadi, Samuel A Lambert, Ishminder Mann, Kate Cook, et al. Determination and inference of eukaryotic transcription factor sequence specificity. *Cell*, 158(6):1431–1443, 2014.

[19] Eugene V Koonin, Yuri I Wolf, and Georgy P Karev. The structure of the protein universe and genome evolution. *Nature*, 420(6912):218–223, 2002.

[20] Michael J Stanhope, Andrei Lupas, Michael J Italia, Kristin K Koretke, Craig Volker, and James R Brown. Phylogenetic analyses do not support horizontal gene transfers from bacteria to vertebrates. *Nature*, 411(6840):940–944, 2001.

[21] Alastair Crisp, Chiara Boschetti, Malcolm Perry, Alan

Tunnacliffe, and Gos Micklem. Expression of multiple horizontally acquired genes is a hallmark of both vertebrate and invertebrate genomes. *Genome Biol.*, 16(50):10–1186, 2015.

[22] Ernesto Pérez-Rueda and Julio Collado-Vides. The repertoire of dna-binding transcriptional regulators in escherichia coli k-12. *Nucleic Acids Res.*, 28(8):1838–1847, 2000.

[23] L. Aravind and E. V. Koonin. Dna-binding proteins and evolution of transcription regulation in the archaea. *Nucleic Acids Res.*, 27(23):4658–4670, Dec 1999.

[24] J L Riechmann, J Heard, G Martin, L Reuber, C-Z Jiang, J Keddie, L Adam, O Pineda, OJ Ratcliffe, RR Samaha, et al. Arabidopsis transcription factors: genome-wide comparative analysis among eukaryotes. *Science*, 290(5499):2105–2110, 2000.

[25] Valérie Ledent and Michel Vervoort. The basic helix-loop-helix protein family: comparative genomics and phylogenetic analysis. *Genome research*, 11(5):754–770, 2001.

[26] Burkhard Morgenstern and William R Atchley. Evolution of bhlh transcription factors: modular evolution by domain shuffling? *Molecular biology and evolution*, 16(12):1654–1663, 1999.

[27] Zeba Wunderlich and Leonid A Mirny. Different gene regulation strategies revealed by analysis of binding motifs. *Trends Genet.*, 25(10):434–440, 2009.

[28] Varodom Charoensawan, Derek Wilson, and Sarah A Teichmann. Lineage-specific expansion of dna-binding transcription factor families. *Trends Genet.*, 26(9):388–393, 2010.

[29] Claire Larroux, Graham N Luke, Peter Koopman, Daniel S Rokhsar, Sebastian M Shimeld, and Bernard M Degnan. Genesis and expansion of metazoan transcription factor gene classes. *Molecular biology and evolution*, 25(5):980–96, may 2008.

[30] Bernard M Degnan, Michel Vervoort, Claire Larroux, and Gemma S Richards. Early evolution of metazoan transcription factors. *Current Opinion in Genetics and Development*, 19(6):591–9, dec 2009.

[31] Mansi Srivastava, Oleg Simakov, Jarrod Chapman, Bryony Fahey, Marie E A Gauthier, Therese Mitros, Gemma S Richards, Cecilia Conaco, Michael Dacre, Uffe Hellsten, Claire Larroux, Nicholas H Putnam, Mario Stanke, Maja Adamska, Aaron Darling, Sandie M Degnan, Todd H Oakley, David C Plachetzki, Yufeng Zhai, Marcin Adamski, Andrew Calcino, Scott F Cummins, David M Goodstein, Christina Harris, Daniel J Jackson, Sally P Leys, Shengqiang Shu, Ben J Woodcroft, Michel Vervoort, Kenneth S Kosik, Gerard Manning, Bernard M Degnan, and Daniel S Rokhsar. The amphimedon queenslandica genome and the evolution of animal complexity. *Nature*, 466(7307):720–6, aug 2010.

[32] M Madan Babu, Sarah A Teichmann, and L Aravind. Evolutionary dynamics of prokaryotic transcriptional regulatory networks. *J. Mol. Biol.*, 358(2):614–633, 2006.

[33] Nacho Molina and Erik van Nimwegen. Scaling laws in functional genome content across prokaryotic clades and lifestyles. *Trends Genet.*, 25(6):243–247, 2009.

[34] Varodom Charoensawan, Derek Wilson, and Sarah A Teichmann. Genomic repertoires of dna-binding transcription factors across the tree of life. *Nucleic Acids Res.*, 38(21):7364–7377, 2010.

[35] Tin Yau Pang and Sergei Maslov. A toolbox model of evolution of metabolic pathways on networks of arbitrary topology. *PLoS Comput. Biol.*, 7(5):e1001137, May 2011.

[36] Marco Cosentino Lagomarsino, Alessandro L Sellerio, Philip D Heijning, and Bruno Bassetti. Universal features in the genome-level evolution of protein domains. *Genome Biol.*, 10(1):R12, jan 2009.

[37] Nacho Molina and Erik van Nimwegen. The evolution of domain-content in bacterial genomes. *Biol Direct*, 3:51, 2008.

[38] Sergei Maslov, Sandeep Krishna, Tin Yau Pang, and Kim Sneppen. Toolbox model of evolution of prokaryotic metabolic networks and their regulation. *Proc. Natl. Acad. Sci. U. S. A.*, 106(24):9743–8, jun 2009.

[39] Shalev Itzkovitz, Tsvi Tlusty, and Uri Alon. Coding limits on the number of transcription factors. *BMC genomics*, 7(1):239, 2006.

[40] Georgy P Karev, Yuri I Wolf, Andrey Y Rzhetsky, Faina S Berezovskaya, and Eugene V Koonin. Birth and death of protein domains: a simple model of evolution explains power law behavior. *BMC Evol. Biol.*, 2(1):18, 2002.

[41] Arttu Jolma, Jian Yan, Thomas Whitington, Jarkko Toivonen, Kazuhiro R Nitta, Pasi Rastas, Ekaterina Morgunova, Martin Enge, Mikko Taipale, Gonghong Wei, et al. Dna-binding specificities of human transcription factors. *Cell*, 152(1):327–339, 2013.

[42] Marcus B. Noyes, Ryan G. Christensen, Atsuya Wakabayashi, Gary D. Stormo, Michael H. Brodsky, and Scot A. Wolfe. Analysis of homeodomain specificities allows the family-wide prediction of preferred recognition sites. *Cell*, 133(7):1277–1289, Jun 2008.

[43] Michael F Berger, Gwenael Badis, Andrew R Gehrke, Shaheynoor Talukder, Anthony A Philippakis, Lourdes Peña-Castillo, Trevis M Alleyne, Sanie Mnaimneh, Olga B Botvinnik, Esther T Chan, Faiqua Khalid, Wen Zhang, Daniel Newburger, Savina A Jaeger, Quaid D Morris, Martha L Bulyk, and Timothy R Hughes. Variation in homeodomain dna binding revealed by high-resolution analysis of sequence preferences. *Cell*, 133(7):1266–76, jun 2008.

[44] Metewo Selase Enuameh, Yuna Asriyan, Adam Richards, Ryan G Christensen, Victoria L Hall, Majid Kazemian, Cong Zhu, Hannah Pham, Qiong Cheng, Charles Blatti, Jessie A Brasefield, Matthew D Basciotta, Jianhong Ou, Joseph C McNulty, Lihua J Zhu, Susan E Celniker, Saurabh Sinha, Gary D Stormo, Michael H Brodsky, and Scot A Wolfe. Global analysis of drosophila cys-his zinc finger proteins reveals a multitude of novel recognition motifs and binding determinants. *Genome research*, 23(6):928–40, jun 2013.

[45] Artem S Novozhilov, Georgy P Karev, and Eugene V Koonin. Biological applications of the theory of birth-and-death processes. *Briefings in bioinformatics*, 7(1):70–85, 2006.

[46] Trevor Fenner, Mark Levene, and George Loizou. A stochastic evolutionary model exhibiting power-law behaviour with an exponential cutoff. *Physica A: Statistical Mechanics and its Applications*, 355(2):641–656, 2005.

[47] Marco Cosentino Lagomarsino, Alessandro L Sellerio, Philip D Heijning, and Bruno Bassetti. Universal features in the genome-level evolution of protein domains. *Genome biology*, 10(1):1–13, 2009.

[48] Sean B Carroll. Evo-devo and an expanding evolutionary synthesis: a genetic theory of morphological evolution.

Cell, 134(1):25–36, 2008.

[49] Matteo D'Antonio and Francesca D. Ciccarelli. Modification of gene duplicability during the evolution of protein interaction network. *PLoS Comput. Biol.*, 7(4):e1002029, Apr 2011.

[50] Jaime Huerta-Cepas, Damian Szklarczyk, Kristoffer Forslund, Helen Cook, Davide Heller, Mathias C. Walter, Thomas Rattei, Daniel R. Mende, Shinichi Sunagawa, Michael Kuhn, Lars Juhl Jensen, Christian von Mering, and Peer Bork. eggnog 4.5: a hierarchical orthology framework with improved functional annotations for eukaryotic, prokaryotic and viral sequences. *Nucleic Acids Res.*, 44(D1):D286–D293, Jan 2016.

[51] Omer An, Giovanni M. Dall'Olio, Thanos P. Mourikis, and Francesca D. Ciccarelli. Ncg 5.0: updates of a manually curated repository of cancer genes and associated properties from cancer mutational screenings. *Nucleic Acids Res*, 44(D1):D992–D999, Jan 2016.

[52] Anastasios Pavlopoulos and Michael Akam. Hox go omics: insights from drosophila into hox gene targets. *Genome Biol*, 8(3):208, 2007.

[53] So Nakagawa, Stephen S Gisselbrecht, Julia M Rogers, Daniel L Hartl, and Martha L Bulyk. Dna-binding specificity changes in the evolution of forkhead transcription factors. *Proceedings of the National Academy of Sciences*, 110(30):12349–12354, 2013.

[54] Ryan O. Emerson and James H. Thomas. Adaptive evolution in zinc finger transcription factors. *PLoS Genet*, 5(1):e1000325, Jan 2009.

[55] Joseph F Ryan, Patrick M Burton, Maureen E Mazza, Grace K Kwong, James C Mullikin, and John R Finnerty. The cnidarian-bilaterian ancestor possessed at least 56 homeoboxes: evidence from the starlet sea anemone, nematostella vectensis. *Genome biology*, 7(7):R64, jan 2006.

[56] Nicholas H. Putnam, Mansi Srivastava, Uffe Hellsten, Bill Dirks, Jarrod Chapman, Asaf Salamov, Astrid Terry, Harris Shapiro, Erika Lindquist, Vladimir V. Kapitonov, Jerzy Jurka, Grigory Genikhovich, Igor V. Grigoriev, Susan M. Lucas, Robert E. Steele, John R. Finnerty, Ulrich Technau, Mark Q. Martindale, and Daniel S. Rokhsar. Sea anemone genome reveals ancestral eumetazoan gene repertoire and genomic organization. *Science*, 317(5834):86–94, Jul 2007.

[57] Elena Simionato, Valérie Ledent, Gemma Richards, Morgane Thomas-Chollier, Pierre Kerner, David Coornaert, Bernard M Degnan, and Michel Vervoort. Origin and diversification of the basic helix-loop-helix gene family in metazoans: insights from comparative genomics. *BMC Evol. Biol.*, 7(1):33, jan 2007.

[58] Nicole King, M Jody Westbrook, Susan L Young, Alan Kuo, Monika Abedin, Jarrod Chapman, Stephen Fairclough, Uffe Hellsten, Yoh Isogai, Ivica Letunic, Michael Marr, David Pincus, Nicholas Putnam, Antonis Rokas, Kevin J Wright, Richard Zuzow, William Dirks, Matthew Good, David Goodstein, Derek Lemons, Wanqing Li, Jessica B Lyons, Andrea Morris, Scott Nichols, Daniel J Richter, Asaf Salamov, J G I Sequencing, Peer Bork, Wendell A Lim, Gerard Manning, W Todd Miller, William McGinnis, Harris Shapiro, Robert Tjian, Igor V Grigoriev, and Daniel Rokhsar. The genome of the choanoflagellate monosiga brevicollis and the origin of metazoans. *Nature*, 451(7180):783–8, feb 2008.

[59] Alessandro Testori, Livia Caizzi, Santina Cutrupi, Olivier Friard, Michele De Bortoli, Davide Cora', and Michele Caselle. The role of transposable elements in shaping the combinatorial interaction of transcription factors. *BMC genomics*, 13(1):400, jan 2012.

[60] Volker Matys, Olga V Kel-Margoulis, Ellen Fricke, Ines Liebich, Sigrid Land, A Barre-Dirrie, Ingmar Reuter, D Chekmenev, Mathias Krull, Klaus Hornischer, et al. Transfac® and its module transcompel®: transcriptional gene regulation in eukaryotes. *Nucleic Acids Res.*, 34(suppl 1):D108–D110, 2006.

[61] Anthony Mathelier, Xiaobei Zhao, Allen W. Zhang, François Parcy, Rebecca Worsley-Hunt, David J. Arenillas, Sorana Buchman, Chih-yu Chen, Alice Chou, Hans Ienasescu, Jonathan Lim, Casper Shyr, Ge Tan, Michelle Zhou, Boris Lenhard, Albin Sandelin, and Wyeth W. Wasserman. Jaspar 2014: an extensively expanded and updated open-access database of transcription factor binding profiles. *Nucleic Acids Res.*, 42(Database issue):D142–D147, Jan 2014.

[62] Arttu Jolma, Teemu Kivioja, Jarkko Toivonen, Lu Cheng, Gonghong Wei, Martin Enge, Mikko Taipale, Juan M Vaquerizas, Jian Yan, Mikko J Sillanpää, et al. Multiplexed massively parallel selex for characterization of human transcription factor binding specificities. *Genome research*, 20(6):861–873, 2010.

[63] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

[64] Huaiyu Mi, Anushya Muruganujan, and Paul D Thomas. Panther in 2013: modeling the evolution of gene function, and other gene attributes, in the context of phylogenetic trees. *Nucleic Acids Res.*, 41(D1):D377–D386, 2013.